MiniMind-3V vision LoRA adapter

Lucas van de weerd | Lucasgitguy-co | 7 November 2026 | Version 1.0

Repository: Lucasgitguy-co/minimind-3v-moe-adapter

Reviewed revision: bca9243cbf6b393d56c0176185773b58750cf8ba

Abstract

This report documents a parameter-efficient vision-language adapter and the evidence needed to load and interpret it. The inspected release supports configuration-level statements, while its historical training procedure and task quality remain incompletely documented.

1. Scope and related work

This technical report reconstructs the release from pinned repository artifacts and the publisher's earlier portfolio report. It is an AI-assisted documentation synthesis, not an independent rerun or a peer-reviewed publication. Reported measurements belong to the original release records. [1, 2, 3]

Low-rank adaptation is described by Hu et al. (2021), LoRA: Low-Rank Adaptation of Large Language Models, https://arxiv.org/abs/2106.09685. This methodological reference does not supply missing results for this adapter.

2. Model and method

The PEFT configuration records rank 8, alpha 16, dropout 0.05 and PEFT version 0.21.2. The declared implementation is MiniMindVLM, with vision_proj saved alongside LoRA modules. The exact historical base checkpoint is unverified. The name alone does not establish a verified mixture-of-experts training run. The base model must be initialized before applying the adapter; these tensors do not form a standalone model. [1, 2, 3]

3. Evidence and results

No documented held-out benchmark or complete run manifest accompanies the inspected adapter. The rank and dropout values describe the saved adapter configuration; they do not reconstruct its training data, number of optimizer updates or measured quality. [1, 2, 3]

4. Limitations and intended use

Reproduction requires the historical base, compatible processor and implementation, exact dataset split, optimizer settings and validation record. Recovering the original base checkpoint is a prerequisite for a faithful reconstruction. The release retains Apache-2.0 attribution.

5. Reproducibility and conclusion

The source snapshot above fixes what was reviewed. The linked source files and source-manifest.json record the evidence used, including file hashes. Consult the repository's usage instructions and recorded dependencies for execution. No new model training, inference benchmark, or claimed state-of-the-art result was introduced in preparing this report. Evidence status: artifact documentation.

References

[1] Lucasgitguy-co. README.md at reviewed revision.

[2] Lucasgitguy-co. adapter_config.json at reviewed revision.

[3] Lucasgitguy-co. docs/PROVENANCE.md at reviewed revision.

[4] Lucasgitguy-co. Local Model Adaptation Across Vision Diagrams Geometry and Game Decisions. Historical portfolio report, 5 November 2026; its proposed v2 experiment is superseded by later release records.